DOCUMENT RESUME

ED 289 477 IR 012 944

AUTHOR Smoliar, Stephen W.
TITLE Review of "Conceptual Structures: Information Processing in Mind and Machine."
INSTITUTION University of Southern California, Marina del Rey. Information Sciences Inst.
SPONS AGENCY Advanced Research Projects Agency (DOD), Washington, D.C.
REPORT NO ISI/RS-87-182
PUB DATE Mar 87
NOTE 11p.; ISI Reprint Series. Reprinted from Artificial Intelligence magazine, 1987.
PUB TYPE Book/Product Reviews (072)

EDRS PRICE MF01/PC01 Plus Postage.
DESCRIPTORS *Artificial Intelligence; Book Reviews; Database Management Systems; *Man Machine Systems
IDENTIFIERS *Conceptual Graphs; *Knowledge Representation; Relational Approach (Information Retrieval)

ABSTRACT

This review of the book, "Conceptual Structures: Information Processing in Mind and Machine," by John F. Sowa, argues that anyone who plans to get involved with issues of knowledge representation should have at least a passing acquaintance with Sowa's conceptual graphs for a database interface. (Used to model the underlying semantics of a database, these graphs describe data from the user's point of view and access data according to the system's point of view. They also represent functional dependencies in a relational database and support inferences and computations that are not explicit in the initial query.) Sowa's ideas on knowledge representation, cognitive science, and expert systems as presented in the book are discussed, and it is argued that: (1) Sowa's ideas were limited by his concern for Codd's relational model of databases; (2) developments in cognitive science during the eight years that elapsed between Sowa's article in the IBM Journal of Research and Development (July 1986) and the publication of the book have raised issues that he does not confront adequately; and (3) perhaps the most valuable legacy of the book will be its impressive set of references to early scholarship. (RP)

****
* Reproductions supplied by EDRS are the best that can be made *
* from the original document. *
****

ED289477

U.S. DEPARTMENT OF EDUCATION
Office of Educational Research and Improvement
EDUCATIONAL RESOURCES INFORMATION CENTER (ERIC)

☒ This document has been reproduced as received from the person or organization originating it.

☐ Minor changes have been made to improve reproduction quality.

• Points of view or opinions stated in this document do not necessarily represent official OERI position or policy.

ISI Reprint Series
ISI/RS-87-182
March 1987

University of Southern California

Stephen W. Smoliar

Review of *Conceptual Structures: Information Processing in Mind and Machine*

Reprinted from
Artificial Intelligence magazine, 1987.

IR012944

INFORMATION SCIENCES INSTITUTE

213/822-1511
4676 Admiralty Way/Marina del Rey/California 90292-6695

BEST COPY AVAILABLE

REPORT DOCUMENTATION PAGE

1a. REPORT SECURITY CLASSIFICATION Unclassified	1b. RESTRICTIVE MARKINGS
2a. SECURITY CLASSIFICATION AUTHORITY 2b. DECLASSIFICATION / DOWNGRADING SCHEDULE	3. DISTRIBUTION / AVAILABILITY OF REPORT This document is approved for public release; distribution is unlimited.
4. PERFORMING ORGANIZATION REPORT NUMBER(S) ISI/RS-87-182	5. MONITORING ORGANIZATION REPORT NUMBER(S) ---
6a. NAME OF PERFORMING ORGANIZATION USC/Information Sciences Institute 6b. OFFICE SYMBOL (If applicable)	7a. NAME OF MONITORING ORGANIZATION ---
6c. ADDRESS (City, State, and ZIP Code) 4676 Admiralty Way Marina del Rey, CA 90292	7b. ADDRESS (City, State, and ZIP Code) ---
8a. NAME OF FUNDING / SPONSORING ORGANIZATION Advanced Research Projects Agency 8b. OFFICE SYMBOL (If applicable)	9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER MDA903 81 C 0335
8c. ADDRESS (City, State, and ZIP Code) 1400 Wilson Boulevard Arlington, VA 22209	10. SOURCE OF FUNDING NUMBERS PROGRAM ELEMENT NO. --- / PROJECT NO. --- / TASK NO. --- / WORK UNIT ACCESSION NO. ---
11. TITLE (Include Security Classification) Review of *Conceptual Structures: Information Processing in Mind and Machine* [Unclassified]	
12. PERSONAL AUTHOR(S) Stephen W. Smoliar	
13a. TYPE OF REPORT Research Report 13b. TIME COVERED FROM _____ TO _____	14. DATE OF REPORT (Year, Month, Day) 1987, March 15. PAGE COUNT 10
16. SUPPLEMENTARY NOTATION Reprinted from Artificial Intelligence magazine, 1987.	

17. COSATI CODES		
FIELD	GROUP	SUB-GROUP
09	02	

18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number)
knowledge representation, conceptual graph, artificial intelligence

19. ABSTRACT (Continue on reverse if necessary and identify by block number)

This is a review of the book *Conceptual Structures: Information Processing in Mind and Machine*, by John F. Sowa, published in 1984 by the Addison-Wesley Publishing Company.

20. DISTRIBUTION / AVAILABILITY OF ABSTRACT ☑ UNCLASSIFIED/UNLIMITED ☑ SAME AS RPT. ☐ DTIC USERS	21. ABSTRACT SECURITY CLASSIFICATION Unclassified
22a. NAME OF RESPONSIBLE INDIVIDUAL Sheila Coyazo Victor Brown	22b. TELEPHONE (Include Area Code) 213-822-1511 22c. OFFICE SYMBOL

DD FORM 1473, 84 MAR 83 APR edition may be used until exhausted. All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE
Unclassified

ISI Reprint Series
ISI/RS-87-182
March 1987

University of Southern California

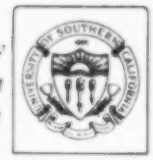

Stephen W. Smoliar

Review of *Conceptual Structures: Information Processing in Mind and Machine*

Reprinted from
Artificial Intelligence magazine, 1987.

INFORMATION SCIENCES INSTITUTE

213/822-1511
4676 Admiralty Way/Marina del Rey/California 90292-6695

This research is supported by the Defense Advanced Research Projects Agency under Contract No. MDA903 81 C 0335. Views and conclusions contained in this report are the author's and should not be interpreted as representing the official opinion or policy of DARPA, the U.S. Government, or any person or agency connected with them.

ISI Reprint Series

This report is one in a series of reprints of articles and papers written by ISI research staff and published in professional journals and conference proceedings. For a complete list of ISI reports, write to

Document Distribution
USC/Information Sciences Institute
4676 Admiralty Way
Marina del Rey, CA 90292-6695
USA

Table of Contents

1. Publication: The Scourge of Timeliness 1
2. Contributions 2
    2.1. Knowledge Representation 2
    2.2. Cognitive Science 4
    2.3. Expert Systems 4
    2.4. Education 4
    2.5. Reading 5
3. Conclusions 5

1. Publication: The Scourge of Timeliness

The July 1976 issue of the *IBM Journal of Research and Development* contained a fascinating article which, at the time of its appearance, seems to have had little impact on the artificial intelligence community (perhaps because of the medium in which it appeared). The article was entitled "Conceptual Graphs for a Data Base Interface;" and the author was John F. Sowa of the IBM Systems Research Institute. The abstract for this paper was brief and to the point:

> A data base that supports natural language queries is not really natural if it requires the user to know how the data are represented. This paper defines a formalism, called conceptual graphs, that can describe data according to the user's view and access data according to the system's view. In addition, the graphs can represent functional dependencies in the data base and support inferences and computations that are not explicit in the initial query.

In other words, Sowa had undertaken to use a graph structure to model the underlying semantics of a data base. The objective of this semantic model was to provide a bridge between "what the user wanted to know" and "what the data base could provide." The graph was structured in such a way that its nodes represented "concepts" and its edges were labeled by "conceptual relations." Sowa's intent was to use these graphs to represent both the concepts modeled by a data base schema and the concepts of interest to a user. He then developed a set of rules for the manipulation of these graphs through which one could transform a user's query, represented as a conceptual graph, into a valid data base query, similarly represented. The example he presented required a modest amount of inferencing and served to illustrate most of the machinery he had discussed.

The paper was a model of scientific publication. It pursued a well-defined course of reasoning in a straightforward manner. The writing style was clean and direct; and the reader who was curious to learn more was informed in a footnote that a book on the subject, entitled *Conceptual Structures: Information Processing in Mind and Machine*, was to be published.

Unfortunately, at about the time this paper was being readied for publication, Ronald J. Brachman was delivering a paper at the Sixth International Conference on Computational Linguistics which would lead one to question just how far one could go with the elegant simplicity of Sowa's foundation. This paper was subsequently published in the *International Journal of Man-Machine Studies* in 1977 under the title, "What's in a Concept: Structural Foundations for Semantic Networks." If Sowa had seen the abstract of this paper, he might have thought it had been directed at him, personally:

> This paper examines the fundamentals of network notation, in order to understand why the "formalism" has not been the panacea it was once hoped to be. We focus here on "concepts"--what net-authors think they are, and how network nodes might represent them. The simplistic view of concept nodes as representing extensional sets is examined, and found wanting in several respects.

In all fairness, Sowa was well aware of the distinction between intensional and extensional representations. However, his view was somewhat confined by his concern for Edward Codd's relational model of data bases. Within this rather narrow universe, he could regard relations as having *intensional* meaning whose *extensions* were the actual tuples in the data base; and ultimately his approach depended on his ability to pass between these two perspectives using the tightrope of Codd's formalism. Brachman, on the other hand, was concerned with broader questions of semantics; and he was discovering that the world was a very complicated place. Subsequent research which has led him to the development of languages like KL-ONE and KRYPTON may be seen to have reinforced his convictions.

It is unfortunate that, to judge from their bibliographies, Sowa and Brachman were unaware of each other's activities. What is more unfortunate is that Sowa's book, which was announced as early as 1976, did not see the light of day until 1984; and what is perhaps *most* unfortunate is that the book which appeared in 1984 is probably not that different in content from the book which was promised in 1976. To be sure, there are entries in the book's bibliography which appeared during this interim period; but

these entries seem to serve more as afterthoughts--weak attempts to stay up to date--than as supports for the body of the text.

Had this book been published in 1976 it would have deserved much attention. In progressing from his paper to his book, Sowa had extended his concern from relational data bases to the workings of the mind. This was a bold leap, but it appears that Sowa took great trouble to collect considerable evidence before taking it. His scholarly sources are impressive, to say the least. If he has read everything he cites, then that accomplishment alone would be creditable. The results of this scholarship provide insights into the domain of cognitive science which are both provocative and intriguing. Such material could have shaken a lot of trees in 1976.

However, while Sowa was busy getting his book into shape, cognitive science was emerging as a significant discipline, burgeoning with practitioners and plans of action. Between 1976 and 1984 ideas which began as piques of the imagination were translated into research programs with well-defined experimental procedures and often fascinating results. One might say that what Sowa had in 1976 was an *hors d'oeuvre* to whet the appetite for cognitive science; but the scientific attitude has responded with gusto over the subsequent eight years, and now one is more concerned with carving the roast. Today's reader will no longer be content to be dazzled by Sowa's observations; he will want to know if anything concrete has emerged from them. On this score, unfortunately, Sowa's text is still essentially back in 1976, so that its most valuable legacy will probably be its impressive set of references to early scholarship.

2. Contributions

It is worthwhile to review the contributions which this book has attempted to make and to assess its success for each of them. Five such possibilities will be considered:

1. Knowledge representation: This is a survey of the issues of knowledge representation and a discussion of how they may be confronted.

2. Cognitive science: This is an attempt to unify the many different views which constitute cognitive science into an overall framework.

3. Expert systems: This is a study of the principles behind expert systems (as opposed to a chronicle of their achievements).

4. Education: This is a textbook for a student who wants to become acquainted with the major issues of artificial intelligence and cognitive science.

5. Reading: This is simply a good book for the reading--one man offering his ideas to his peers.

2.1. Knowledge Representation

As was just observed, the knowledge representation pot was already coming to a boil in 1976. Unfortunately, Sowa's acknowledgment of this situation, or of the questions which were being raised by it, never rises above the passing reference. Consequently, the issues of knowledge representation are addressed only to the extent that Sowa can introduce devices with his conceptual graphs which handle them. (Anyone who has seen *A Chorus Line* will probably remember a show-stopping number called "I Can Do That.") After a while, the devices are flying so fast and furious that one has to catch one's breath to remember what the underlying issues were. This is the fatal moment, because it is at this point that one realizes that Sowa has not been terribly clear about what his concepts really are or how they are to be used. When he was dealing with relational data bases, he was able to be more direct: concepts essentially corresponded to the domains of a relational data model, while a single conceptual graph served to model the intensional semantics of a relation. However, to attempt to generalize this approach to a broader view of knowledge amounts to an attempt to develop a theory of epistemology based on the

relational data model; and while this is an approach which *has* been seriously investigated, that investigation has raised issues which Sowa's book does not confront adequately.

The reader who is in search of a more substantive discussion of such issues would do better to seek out a copy of Brian Smith's Ph. D. Dissertation of 1982 at the Massachusetts Institute of Technology, "Reflection and Semantics in a Procedural Language." Sowa's bibliography does not include this document, so it is not surprising that Sowa has omitted any reference to what Smith came to call the ***knowledge representation hypothesis***:

> Any mechanically embodied intelligent process will be comprised of structural ingredients that a) we as external observers naturally take to represent a propositional account of the knowledge that the overall process exhibits, and b) independent of such external semantical attribution, play a formal but causal and essential role in engendering the behavior that manifests that knowledge.

From Smith's point of view, one might say that Sowa has tended towards putting all his eggs into the first of these two baskets: developing a propositional account of all external observations. However, it remains unclear what role is to be played by the structures which support that *propositional* account when it comes to a ***behavioral*** account of knowledge.

An example of the imbalance between propositional and behavioral accounts may be found in Sowa's discussion of logical deduction. For Sowa, the way one handles deduction is to introduce a `PROPOSITION` type. Having posited the type, one can then apply logical rules of inference to instances of it. The fact that the behavioral issue of determining *how* to perform such inferences remains a problem of imposing magnitude, regardless of the representation of the propositions, doesn't seem to faze Sowa. All that seems to be important is that it should be possible to draw a graph (i.e. give a propositional account) for any aspect of knowledge which the intuition (i.e. external observation) can grasp.

Nevertheless, even the intuition may wander beyond the modeling abilities of conceptual graphs. Thus, when the discussion turns to operators, some of the graphs appear more like dataflow diagrams. New shaped nodes are introduced to represent processing elements. In other words, a new propositional syntax has been imposed. Unfortunately, it is never made terribly clear just how it relates to Sowa's initial notation, except for the fact that the entities which are manipulated by the processing elements are supposed to be instances of types.

Indeed, Sowa shows a general tendency to get his syntactic act together and take it on the road before considering whether or not he should bring along some semantics as well. While his book has a section entitled "Model Theory," his exposition of this subject tends to dwell on a view of a model as an account for the contents of a relational data base. This approach is naive, to say the least, when one considers the grander plans that Sowa has for his conceptual structures; and Alfred Tarski, who first raised many of the major issues of model theory, is rather summarily dismissed in a comparison of the distinction between model theory and proof theory with the distinction between procedural and declarative representations of knowledge.

Ultimately, it would appear that Sowa is trying to address questions of knowledge representation by drawing pictures. However, his approach breaks down for two major reasons. First of all, he tends to evade any commitment to a set of questions to be addressed. In retrospect, one may say that, *de facto*, he has concerned himself with a propositional account of knowledge; and it is possible that, due to his lack of familiarity with the knowledge representation hypothesis, the need for a corresponding behavioral account may never have occurred to him. The second reason is that pictures do not answer questions. Their power lies in their ability to provide useful ***abstractions*** which one may then use as tools to answer questions. Thus, Sowa's approach to knowledge representation is the reverse of a conventional scientific inquiry: The reader must infer the questions being addressed from the abstractions implicit in the pictures, rather than proceed along the more traditional path from questions to useful abstractions to graphic illustration.

2.2. Cognitive Science

Sowa has a lot to say about the brain in terms of both its hardware and its software. Furthermore, most of what he has to say is well-cited. Since I am not, by profession, a brain scientist, I cannot evaluate the timeliness of this material as I can the material for knowledge representation. I can, however, comment that I found most of Sowa's presentations of this material to be uncomfortably categorical. I felt as if I were being given a gospel rather than being led down a path of inference and discovery. I acknowledge that this is a purely subjective response (which is why I have lapsed into the vertical pronoun). Suffice it to say that while Sowa can provide the reader with the fruits of his literature search, he does not appear to have been engaged in active research in this area. Consequently, there are other writers with a more substantive track record on this subject who do a much better job of instilling confidence in their readers with regard to both the points they are trying to make and their use of the available literature for the support of those points.

Where Sowa is clearly at his weakest is on the topic of the modeling of human memory. Although Roger Schank's *Dynamic Memory* is included in the bibliography, there is no reference to it in the body of the text. While Sowa *is* aware of some of Schank's work on memory organization packets (MOPs), this is discussed in the chapter on "Reasoning and Computation." Thus, MOPs are presented as a device for modularity, rather than a model of memory. Similarly, Janet Kolodner's application of the MOP principles to CYRUS, which attempted to address such critical issues as forgetting, are totally ignored. CYRUS is cited merely for its contribution (with FRUMP) to "a typical example of an AI database."

2.3. Expert Systems

The contribution here is extremely weak. It would appear that when this book was first being promised in 1976, the core of this portion would be concerned with a discussion of the nature of the reasoning performed by the system described in Sowa's paper. Unfortunately, this discussion has little to do with the modeling of expert knowledge. Indeed, the issue of *expertise* as related to either the design or the use of data bases is never addressed. Thus, Sowa is not so much interested in the "expert" nature of expert systems as he is in the potential for "applied artificial intelligence." Ultimately, this portion of the book provides a bevy of citations which fly the reader past the topics of natural language interfaces, knowledge acquisition, learning, and examples of expert systems. None of these topics is discussed in any significant depth, and the curious reader is left with few clues as to which of those citations should be pursued for more intensive study.

2.4. Education

One of the "missions" of the IBM Systems Research Institute is educational. One may then conjecture that this book is the product of a course which Sowa has taught there. Thus, its initial intended role was that of a textbook. Such a conjecture is reinforced by the abundance of exercises and the suggestions for further reading at the end of each chapter.

Textbooks for artificial intelligence courses are hot items. The area is churning with competition and strong opinions. Fortunately, this is not the place to attempt to survey the possibilities facing any potential teacher of this subject. However, it should be clear from the above comments that this book is *not* particularly appropriate for a course being given in the present day. For all the appearance of breadth, its scope is ultimately too narrow. Most damaging, however, is the problem that the categorical approach to many issues will be more likely to mislead than inform. This problem is compounded by the fact that Sowa tends not to distinguish between fact and opinion. An experienced scholar with a well-developed sense of skepticism may be able to use this book as an introduction to a new area of inquiry. Greener students who may still be susceptible to gullibility would do well to steer clear.

2.5. Reading

Every book should be at least considered, although not ultimately judged, for its entertainment value. It is a sad truth that there is quite a lot of bad writing in the world of artificial intelligence, and most of it tends to be reinforced by bad editing. At the very least one should ask of every author how much respect he shows the writer's craft.

This is a highly subjective area for evaluation. Therefore, I must again revert to the first person and state simply that I found the style of this book to be very dull. As a matter of fact, it is testimony to the quality of Sowa's 1976 paper that I read it *after* I read his book, at a time when I was beginning to feel I had had my fill of him. In such a hostile emotional climate the paper shone forth as an object lesson for authors of technical papers. From this I have concluded that Sowa the writer fell victim to the magnitude of his project. His ability to create interest in a topic could be sustained for the length of a paper, but the book fell beyond the limit of his writing skills.

What is missing from this text is a sense of involvement or enthusiasm. Perhaps length is only part of the problem. Time may be another significant element. One may assume that the involvement and enthusiasm *were* there in 1976. However, the processes of delay and rewriting can be erosive. If one lingers over one's own words too long, those words begin to lose their effect and one loses one's passion for them. Then the process is reduced to the routine of grinding out a product, and the effect of that process is not lost on the reader.

3. Conclusions

Anyone who plans to get involved with issues of knowledge representation should have at least a passing acquaintance with Sowa's conceptual graphs. However, one is more likely to come away with a useful outlook of this work from Sowa's paper of 1976 than from his book. Of course, there is much more to knowledge representation than Sowa's view of it. In this respect there is no textbook which will serve the needs of the curious student. However, the *Readings in Knowledge Representation* volume prepared by Ronald J. Brachman and Hector J. Levesque will provide one with most of the current perspectives on the subject. (Indeed, one would hope that if this volume goes into a second edition, then some space may be set aside for Sowa's paper.)

Cognitive science is another area which is not very well served by any textbook. David E. Rumelhart and Donald A. Norman have prepared an excellent chapter on "Representation in Memory" for the second edition of the Steven's *Handbook of Experimental Psychology*; but this volume is not yet published (and one hopes it will not suffer the same fate as Sowa's book). Copies of this chapter may still be available as a technical report from the authors at the Center for Human Information Processing at the University of California in San Diego.

In the area of expert systems, there seems to be no end of books which provoke no end of opinions. However, none of these books have been any more successful than Sowa's in their attempts to provide a unified framework for the subject. One can only hope that eventually such a book will be written.